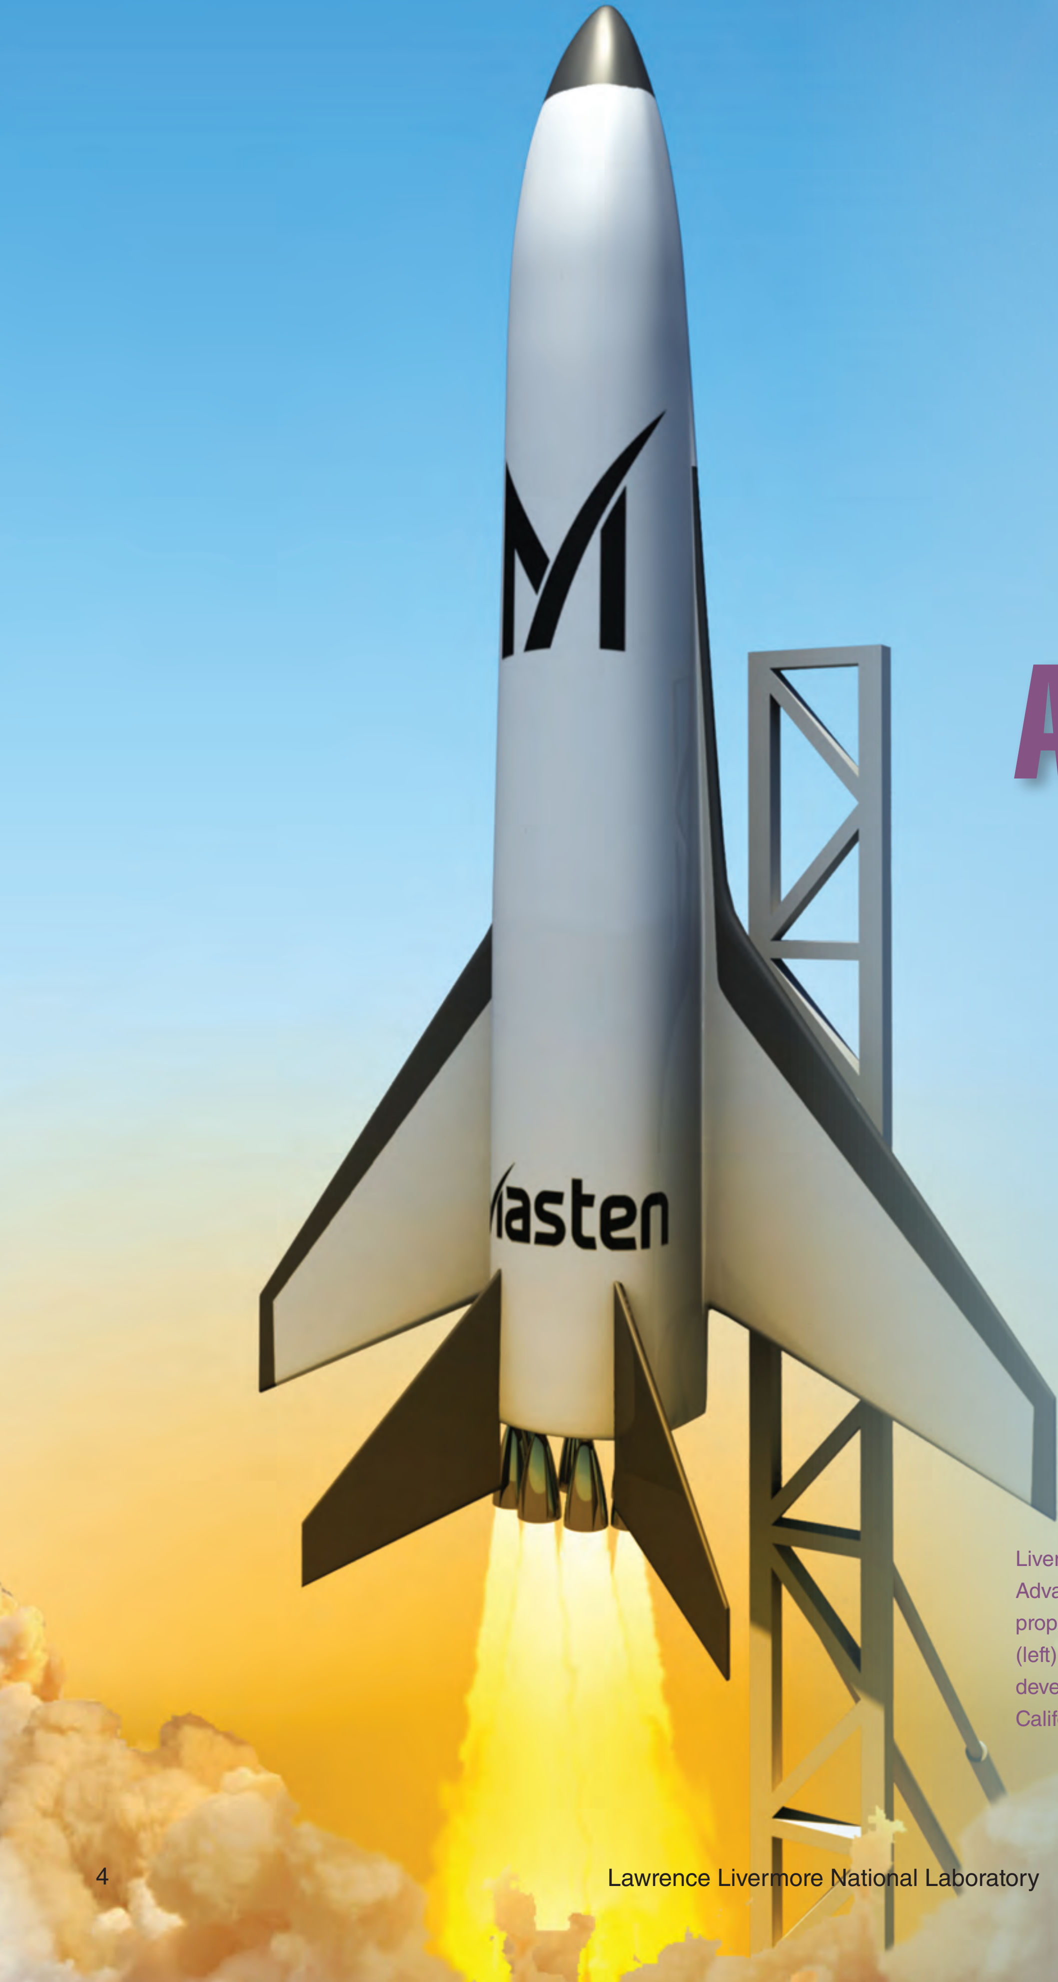

Advancing and the

Livermore researchers demonstrate the power of high-performance computing for developing aerospace technologies.

Livermore engineers are working with the Defense Advanced Research Projects Agency (DARPA) on a proposed medium-lift-capacity XS-1 space plane. (left) An artist's rendering illustrates an early design developed by Masten Space Systems of Mojave, California. (Figure courtesy of Masten Space Systems.)

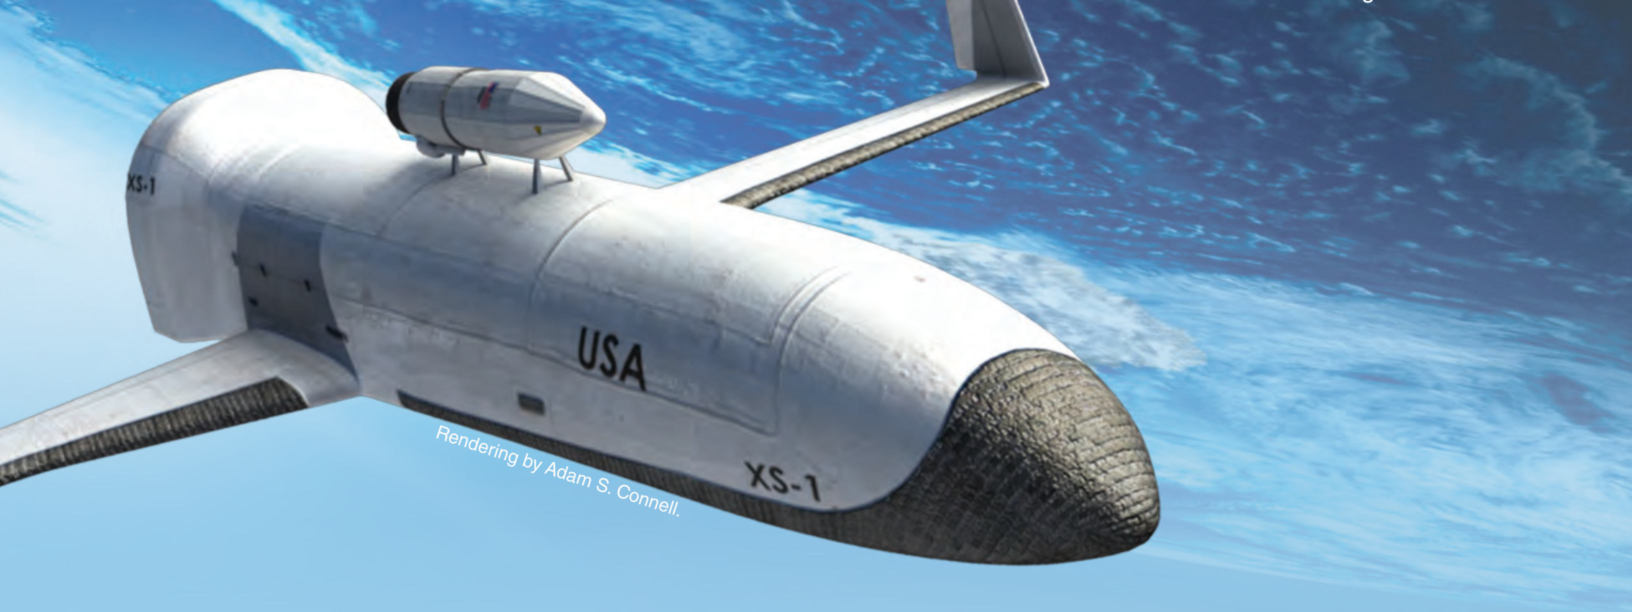

Rendering by Adam S. Connell.

Next-Generation Rockets Engines that Power Them

FOR more than two decades, Lawrence Livermore researchers have pioneered the use of advanced simulation software and some of the world's most powerful supercomputers to virtually design, prototype, and test new components—and entire engineering systems—for national security missions. This approach represents a radical departure from legacy industry practices in which prototypes are typically designed, built, field-tested, broken, and then reworked, significantly increasing the time and expense of evaluating prototype architectures and materials.

In the past few years, Livermore engineers have exploited high-performance computing (HPC) to hone concept designs before prototypes are built and tested. In this way, advanced military systems that work flawlessly are delivered more quickly and at reduced cost. For example, the Livermore-developed BLU-129B low-collateral-damage munition, built for the Department of Defense (DOD), was produced in record time—the first prototype was produced in months rather than the typical years. (See *S&TR*, March 2013, pp. 4–9.) More recently, an interdisciplinary team of Livermore scientists and engineers used HPC to complete a shorter and significantly less expensive development and testing program for a hypersonic conventional warhead. The effort culminated in a sled test proving the design of the warhead and aeroshell and its carbon-based materials. (See *S&TR*, December 2014, pp. 4–11.)

In response to a national need for routine—and more affordable—access to space, Livermore researchers are developing new simulation tools, techniques, and expertise to make possible cost-effective rocket-engine and launch-vehicle designs for national security and scientific exploration. These new capabilities promise to significantly reduce the time, expense, and risk of realizing new space technologies. Livermore engineers and scientists are also creating and applying new additive-manufacturing methods to space technology applications that will enable faster and cheaper routine production of complex parts. (See the box on p. 11.)

Two Efforts for One Agency

In late 2014, a team of Livermore aerospace, mechanical, computational, materials, electronics, and systems engineers began work on two space-technology efforts for the Defense Advanced Research Projects Agency (DARPA). For the first project, researchers conducted a series of simulations directed at evaluating a novel liquid-propellant rocket engine (LRE) design intended for DARPA's Next-Generation Rocket (NGR) program. The simulations focused on the engine's injector performance, combustion characteristics, cooling system, thermal and structural characteristics, and altitude-compensating ability. "DARPA gave us 'seedling' funding to demonstrate our HPC and engineering capabilities for evaluating a new engine design," explains Bill Bruner, Livermore's NASA relationship manager.

The second effort used simulations for studying the thermal and structural response of a candidate launch-vehicle design for DARPA's proposed medium-lift-capacity XS-1 space plane. The XS-1 is envisioned as the first stage of a reusable launch vehicle that could fly 10 times in 10 days, exceed Mach 10 speeds at least once, and launch a payload of up to 2,267 kilograms into low-Earth orbit.

"We're showing DARPA managers we can use computation and modeling to decrease risk, expense, and development time," says aerospace engineer and computational physicist Greg Burton, who leads the Turbulence Analysis and Simulation Center (TASC) within Livermore's Computational Engineering Division. Most of the simulations for the two DARPA projects were conducted under the auspices of TASC.

Engineer Bob Addis, who helped manage Livermore's previous conventional warhead and sled test programs, notes the importance of HPC's cost-savings potential. He explains that in the past it has been more economical for NASA and DOD to purchase liquid-propellant rocket motors from Russia, rather than fund development of U.S.-made motors using legacy industry design practices and techniques. Addis and others are encouraged by the emergence of small, visionary U.S. firms with new ideas for putting payloads reliably and inexpensively into orbit, as well as by DARPA's interest in funding research to develop American-made next-generation launch vehicles and rocket engines.

Computer-Testing a Rocket Engine

The first DARPA-sponsored effort was a partnership between Livermore and WASK Engineering of Cameron Park, California. Livermore engineers were tasked with creating and applying numerical tools to simulate WASK's design of an advanced LRE with an aerospike nozzle. The nozzle design features modular rocket thrust cells arranged in a ring around a novel zero-length central spike.

Measuring about 5 centimeters long and 2 centimeters across, each thrust cell can burn for 5 to 8 minutes and be reused at least 100 times. The cells feature an unusual fuel-injection system for mixing methane (fuel) with liquid oxygen (oxidizer). Each thrust cell resembles an inverted old-fashioned glass milk bottle, with the neck of the bottle representing the combustion chamber's throat. The combustion gases discharge to the nozzle (mouth of the bottle), where the gases are expanded, creating a powerful exhaust. Engineer Allen House, manager for both the aerospike LRE and launch-vehicle simulation projects, notes that WASK provided computer-aided design models of the thrust cell, as well as estimated values for combustion temperatures, pressures, and fuel flow rates.

Aerospikes may be a more efficient alternative to the traditional bell nozzle, in which the exhaust emerges from the combustion chamber and expands against the fixed geometry of the nozzle wall, producing much of the engine's thrust. Bell nozzles operate at peak efficiency when the nozzle expands the exhaust to ambient atmospheric pressure at its exit. However, because ambient atmospheric pressure decreases as a rocket accelerates to orbit, bell-nozzle engines are most efficient at only one particular altitude. As a result, the rocket operates at less than peak efficiency through most of its boost to orbit, and thus must carry more fuel, thereby increasing its weight and decreasing both the altitudes it can reach and the weight of payload it can deliver.

Aerospike LREs, on the other hand, naturally compensate for changes in atmospheric pressure. Rather than expanding gases along a bell nozzle's fixed geometry, the aerospike both expands gases along the fixed wall of the central spike, which provides thrust to lift the vehicle, and to the open atmosphere on the other side. As the

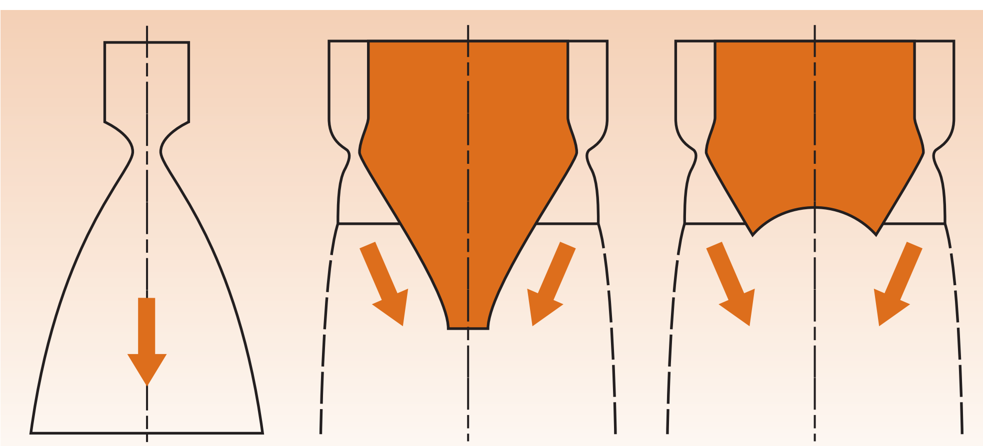

(left) A traditional bell nozzle, (middle) an aerospike nozzle with 50-percent length, and (right) a zero-length aerospike nozzle all produce thrust as the exhaust gases press against the nozzle surface. The two aerospike configurations vent exhaust gases to the ambient environment on the outside of the spike but are able to maintain higher exhaust pressure at the spike surface than a traditional bell nozzle, thereby increasing the engine's thrust and efficiency.

spacecraft climbs to higher altitudes, the air pressure holding the exhaust against the open side naturally decreases, keeping higher pressure on the spike side. In this way, the exhaust geometry automatically adjusts to changes in pressure, and the engine more efficiently uses its available fuel over a wide range of altitudes.

Inside Aerospike Engines

Simulating any LRE is a multiphysics problem that includes transport and mixing of oxidizer and fuel at challenging pressures, densities, and temperatures; complex combustion chemistry; and convective and radiative thermal transport. Coupling and studying these phenomena requires prodigious computational resources. Fortunately, Burton's team had at its disposal the Laboratory's massively parallel supercomputers that use tens of thousands of processors in tandem. These machines were used to run established Livermore codes, commercial codes, and two computational fluid dynamics (CFD) codes that were initially developed at Stanford University's Center for Turbulence Research. Together, these codes describe the flow and mixing of methane and liquid oxygen; reactions and concentrations of the chemical compounds involved in combustion; efficacy of the engine-cooling system; temperature distribution and resultant thermal stresses important for analyzing the structural strength of the thrust cells; and structural and thermal loads during takeoff, flight, and landing.

The two Stanford turbulent mixing and combustion codes, CharlesX and JOE, were originally developed as part of the Department of Energy's (DOE's) Predictive Science Academic Alliance Program to simulate combustion in a high-Mach-number scramjet engine—an air-breathing jet engine in which combustion occurs in a supersonic stream of gas. The codes were substantially modified and enhanced by the Livermore engineering team so they could accurately simulate a functioning LRE like the WASK aerospike design.

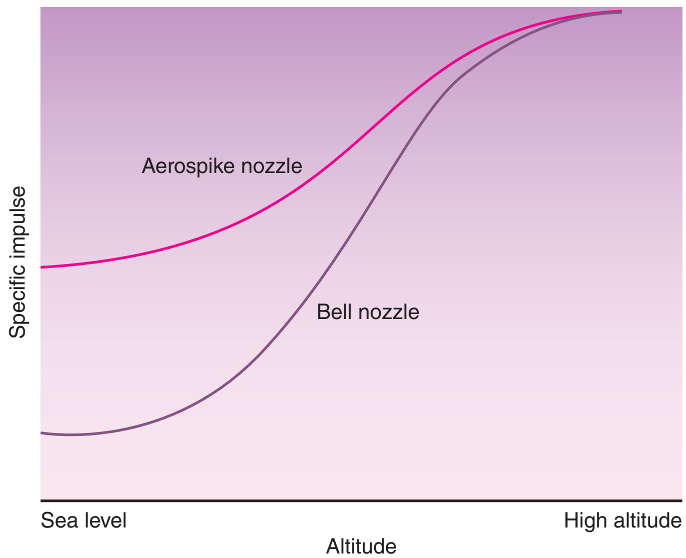

A bell-nozzle engine's peak efficiency is achieved when the nozzle expands the exhaust at its exit to the ambient atmospheric pressure at a particular altitude, and achieves lower efficiency both above and below that altitude. This notional graph shows the efficiency (specific impulse) of an aerospike design versus a bell nozzle optimized for a high altitude. The aerospike achieves more thrust compared to a bell nozzle at most altitudes of operation.

Sister Codes CharlesX and JOE

CharlesX is a state-of-the-art turbulent combustion large-eddy simulation (LES) code for resolving the transient events that occur during turbulent mixing of fuel and oxidizer and their subsequent combustion. The simulations detail important features such as heating and other high-stress transients, which can prematurely age or damage engine walls or potentially cause catastrophic failure of the engine and the launch vehicle. By contrast, sister code JOE is a Reynolds-Averaged Navier Stokes (RANS) simulation code, which resolves only steady-state features.

Simulations using both codes were conducted on Livermore's Cab and Syrah supercomputers. Team members Matt McNenly, Nick Killingsworth, Ryan Vignes, and summer student Elyce Bayat conducted three-dimensional (3D) simulations on two meshes, a moderately resolved 9-million-cell mesh, often used for RANS simulations, and a much finer 60-million-cell mesh for LES runs. The 9-million-cell mesh permitted millimeter-to-nearly-micrometer spatial resolution, thereby allowing simulations to proceed faster than with finer meshes. "It helped us quickly examine coarse features within the system with reasonable accuracy," says Burton. The 60-million-cell mesh resolved certain transient flow and combustion features on the 10-nanosecond and 10-nanometer scale.

The simulations used as many as 4,096 cores (central processing units) over 200 hours of processing time, to successfully isolate features of the transport, mixing, and combustion of the oxidizer and fuel in a thrust cell. The result was one of the highest fidelity 3D simulations of an operating rocket engine ever performed. Burton says the simulations provided insight into how the WASK injector design affects performance, as well as the temperatures and pressures generated during combustion that may affect engine efficiency and stability.

Implementing such fine meshes and large numbers of processors is unprecedented for LRE simulations. Says Bruner, "HPC to traditional rocket designers typically means using one-hundredth the number of cores that we do. Such simulations can provide some design insight but do not allow the level of risk reduction we can achieve with many times additional computing power." Burton

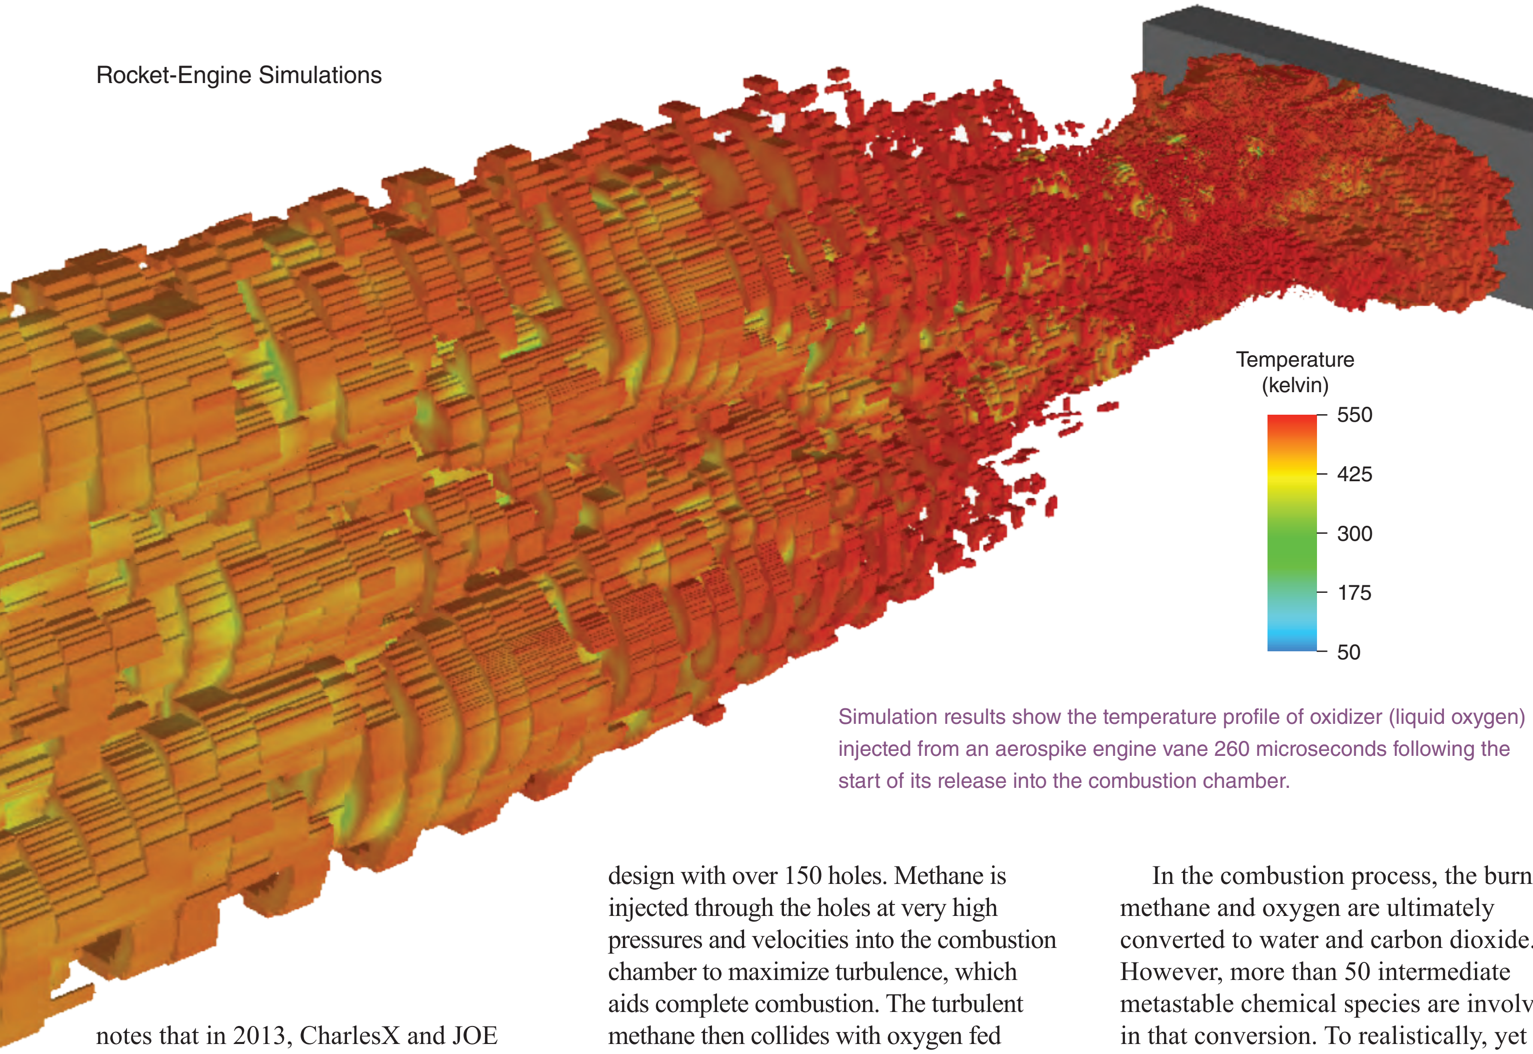

Simulation results show the temperature profile of oxidizer (liquid oxygen) injected from an aerospike engine vane 260 microseconds following the start of its release into the combustion chamber.

notes that in 2013, CharlesX and JOE were successfully run on other projects at high efficiency when scaled up to 1.5 million cores on Livermore's Sequoia supercomputer. Bruner and Burton look forward to future work where simulations of this magnitude may significantly improve accuracy of results.

A Turbulent Process

The team first examined how the system mixes fuel and oxidizer at the head of the combustion chamber. The researchers used a computational mesh that almost exactly reproduced the geometry of the injection system, which features a "showerhead" design with over 150 holes. Methane is injected through the holes at very high pressures and velocities into the combustion chamber to maximize turbulence, which aids complete combustion. The turbulent methane then collides with oxygen fed into the chamber through a series of vanes. Burton's team also performed both RANS and LES simulations of a single vane to better understand how the oxidizer interacts with the turbulent methane. The simulations confirmed previous experimental studies that showed the dominant mixing mechanisms. The results were then used to initialize simulations of the combustion chamber.

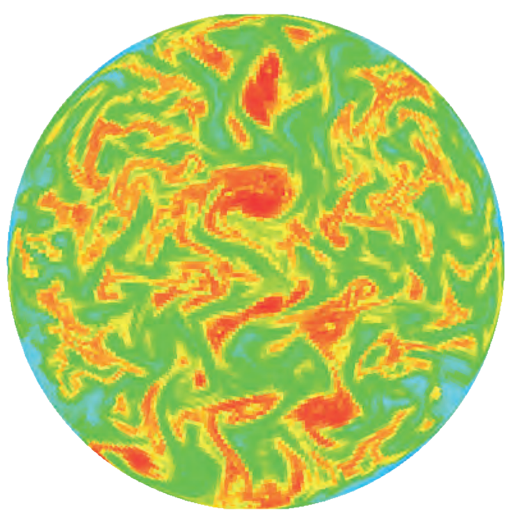

Using CharlesX software and a 9-million-cell mesh, researchers analyzed temperature fields within an aerospike thrust-cell combustion chamber. This cross section of the chamber shows real-time combustion details on a submicrosecond timescale as well as the location of flame fronts (shown in orange).

In the combustion process, the burning methane and oxygen are ultimately converted to water and carbon dioxide. However, more than 50 intermediate metastable chemical species are involved in that conversion. To realistically, yet efficiently, simulate this complicated process, the Livermore team combined CharlesX and JOE with Stanford's Flamelet Progress Variable (FPV) model. The FPV model uses a pre-tabulated chemistry database to relate traditional measures of a combustion system, such as species concentration and heat release, to other variables calculated by the simulation. A combustion–reaction model developed by McNenly was combined with FPV to incorporate these chemical intermediaries and provide a more complete picture of combustion. "Chemicals react and produce intermediaries when they burn," explains McNenly. "This approach is an effective way to capture the most important aspects of the combustion reaction."

By tracking the mixing of fuel and oxidizer and the complex chemical reactions involved in combustion, the simulations show how efficiently and stably the engine can generate the

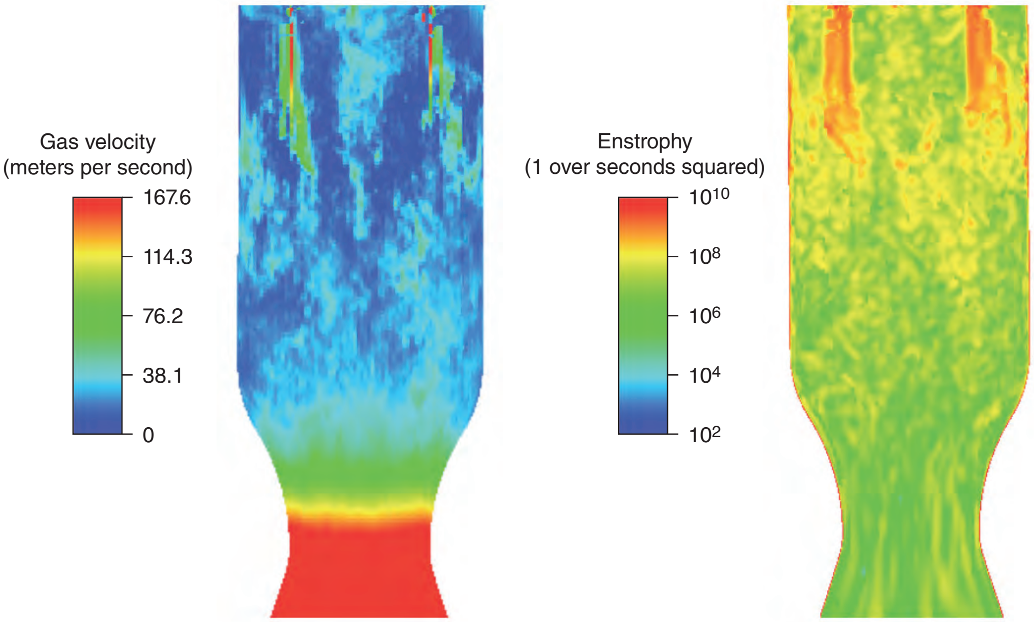

Researchers applied a 9-million-cell mesh to simulate the operation of a thrust-cell combustion chamber. Longitudinal cross sections of the chamber show (left) the gas exhaust velocity and (right) the enstrophy—a measurement that approximately denotes the intensity of the mixing process.

hot gases needed to power a launch vehicle. LES combustion studies yielded instantaneous "snapshot" views of mixing and flame dynamics throughout the engine from the injector and then combustion chamber, to downstream of the throat by the nozzle. The simulations reveal random, tiny, and highly transient phenomena when burning stops (called blowouts), often followed by reignition. Such phenomena can create temperature spikes that make combustion unstable and capable of engine damage. McNenly notes that the simulations indicate that the faster the mix, the more complete the combustion.

Keeping It Cool

An integral part of the aerospike's proper functioning is its cooling system. "The combustion chamber must be continually cooled or it will burn up," explains Livermore engineer Pete Fitsos. Combustion temperatures can reach 3,500 kelvin, but the cooling system must reduce those to 700 kelvin at the thrust cells' outer walls.

The combustion chamber is composed of a high-conductivity copper-alloy liner surrounded by a structural inconel (high-performance alloy) barrel. To maintain uniform cooling and avoid damaging thermal stresses, cryogenic methane flows through a network of 2-millimeter-wide channels that wrap around the copper liner. The warmed methane is then injected into the combustion chamber.

To model the interplay between the cooling liquid and the combustion chamber, Fitsos began with a temperature–pressure profile generated by the CFD simulations. He then used commercial codes to obtain a temperature distribution describing how heat is transferred to an individual cooling channel. The result was a 3D view of the locations and magnitudes of the stresses generated as the copper liner expands and pushes against the inconel frame. The temperature distribution was imported into a model that calculated the overall structural loads on the thrust cell from combustion-generated temperature gradients and pressures.

A reusable rocket engine undergoes punishingly large temperature and pressure swings during its use. Calculating thermal stresses is important because they can be the limiting factor in the lifetime of the engine. Fitsos says data generated by the simulations can indicate areas that might become overstressed during normal operation, which could result in engine and launch-vehicle failure. Follow-up simulations can then determine the effectiveness of any proposed design modifications. Fitsos says, "We have shown that with our computational resources and the right codes, we can do more detailed work than was previously possible."

Simulating External Aerospike Flow

Near the end of the project, Burton agreed to conduct the first-ever 3D LES study of the time-varying turbulent flow from an external aerospike exhaust plume. Starting with the thrust-cell exhaust conditions indicated by the earlier internal combustion simulations, he ran three simulations at altitudes of 3,000; 15,000; and 30,000 meters using the full 3D aerospike geometry. The same grid size was used for each simulation, with the highest resolution aimed at capturing the intense turbulent dynamics generated by the thrust cells' exhaust on the aerospike engine's surfaces. The simulations ran on 2,048 cores of Cab for 250 hours each.

The resulting exhaust flow fields were used to estimate overall engine performance and to evaluate the altitude-compensating behavior of the aerospike engine design. To complete these new

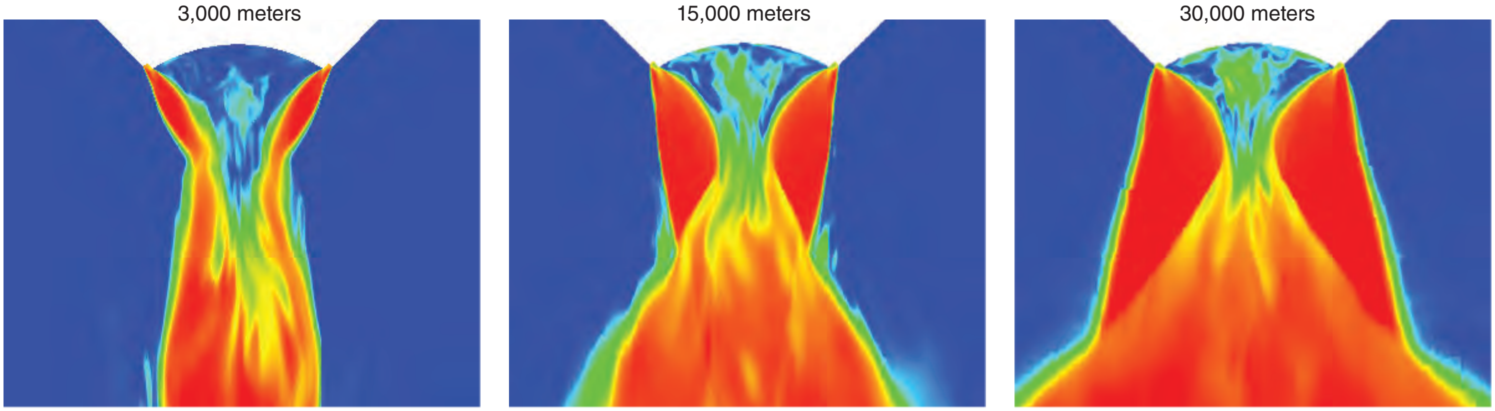

Three simulations show the exhaust flows of an aerospike engine at altitudes of 3,000; 15,000; and 30,000 meters. The simulations reveal that in contrast to conventional rocket-engine designs, the aerospike design compensates for altitude differences.

simulations on schedule and within budget, the team introduced certain approximations that led to specific impulse (a measure of rocket-engine efficiency) estimates nearly 8.5 percent smaller than prior analytical estimates. Burton expects that in future work configured without the approximations, results will range much closer to analytical estimates.

Importantly, "We calculated that the engine does in fact compensate for altitude differences, and in the same amount predicted by analytical studies," says Burton. "We can thus say with some confidence that at least this particular aerospike design exceeds the performance of conventional bell engines."

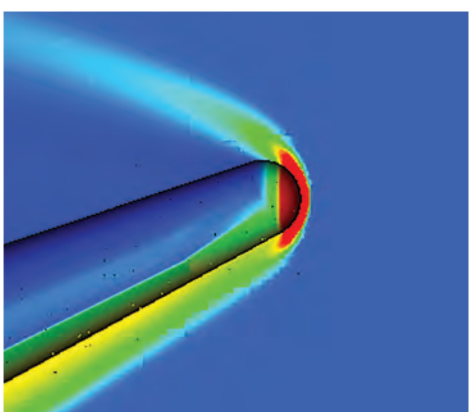

Livermore simulations have shown the temperature distribution on the nose of a hypersonic vehicle, such as the XS-1, as it re-enters the atmosphere. Colors denote relative temperatures.

Assessing Structural Soundness

The aerospike engine could one day fly on a launch vehicle like DARPA's proposed XS-1 hypersonic vehicle. Livermore engineers provided DARPA aerodynamic–thermal–structural analysis on a proposed vertical-takeoff and vertical-landing design concept for the XS-1 developed by Masten Space Systems of Mojave, California. The team combined aerodynamic and thermal stresses generated by Burton's CFD simulations with structural modeling using NIKE3D, a code developed at Livermore. The engineers assessed the mechanical integrity of the Masten design at various points along a simulated flight trajectory from liftoff to Mach 10 and back to vertical landing.

Over several months, Burton and engineer Will Elmer iterated with Masten design engineers to evaluate the structural integrity of the company's evolving XS-1 design. "We aided Masten in selecting structural designs that satisfied DARPA's flight requirements," says Burton. The effort included analyzing critical components such as the vehicle's liquid fuel tank. In one instance, the simulations showed the superiority of a design refinement to structural components supporting the fuel tank.

"Masten needs to be confident that their launch-vehicle design is controllable under all flight conditions," says Elmer. One focus of the study was the performance of the plane's four flaps (or fins) that provide important aerodynamic stability. Elmer explains,"We analyzed the drag and lift on the flaps to ensure no oscillation of the vehicle would occur in flight." The team also examined the integrity of the flaps' connections to the vehicle.

RANS and LES codes analyzed thermal and pressure stresses resulting from the extreme changes in temperature and outside aerodynamic pressures as the launch vehicle fulfilled its mission to place a payload into low-Earth orbit and return to the launch pad. The engineers provided snapshots of the structural and thermal loads on the vehicle at points along its simulated flight trajectory. "Through this work, we can show companies how to catch problems early and save money," says Elmer.

Ready to Partner

"With the past successful DOD programs and the most recent work for DARPA, we have demonstrated what's possible in a short timeframe with HPC," says Bruner. "We can model combustion and flight processes with a level of fidelity that can significantly shorten expensive development cycles, while increasing confidence in the design." As a result, Bruner envisions engaging to a much greater degree with the space flight community to offer Livermore's computational and additive-manufacturing services. However, he

Additively Manufactured Rocket Parts

An efficient technology design cycle requires rapid and cost-effective prototyping. The complex shapes of many rocket components makes additive manufacturing a faster and cheaper tool for providing prototype—as well as flight-ready—parts, when compared to conventional fabrication techniques. For example, the myriad tiny channels of an aerospike thrust-cell cooling system pose significant manufacturing challenges.

Over the last few years, Livermore engineers and materials scientists have demonstrated additive-manufacturing capabilities for producing parts with unique properties pertaining to density, strength, extreme geometry, and surface finish. These capabilities could produce liquid propellant rocket engines (LREs) for components that are stronger, more efficient, lighter, and heat resistant than those available from traditional manufacturing techniques. Such components could improve the efficiency of the engine, allowing larger payloads to be launched with reduced risk into low-Earth orbit.

In 2013, Livermore engineers demonstrated the potential of additive manufacturing when they worked with a small-firm rocket designer. The engineers successfully "printed" a full-scale prototype LRE containing convoluted cooling channels that could not be made by conventional manufacturing techniques. Livermore materials engineer Stephen Burke supervised the effort. Burke started with a computer-aided design file provided by the customer, which was converted to an industry-standard file appropriate for three-dimensional printing. The motor was made from powdered stainless steel particles measuring 30 micrometers in diameter. The particles were fused together using Livermore's laser-welding machine, which laser-welds one layer of powdered particles at a time. The machine ran unattended over 8 days, 24 hours per day, producing thousands of layers, each 25-to-30-micrometers thick.

Burke says that using conventional manufacturing methods, the rocket motor would have had to been made in several pieces and then welded together. "We built one complete piece about 26 centimeters long and 10 centimeters in diameter," he says. "We can build parts now that used to be impossible to manufacture."

Livermore engineers manufactured this rocket motor with an additive-manufacturing process using stainless steel particles measuring 30 micrometers in diameter.

emphasizes that any partnership must be the "right fit," working with a company that understands the enormous potential of HPC.

Burton and his team have been spreading the word about Livermore's space-technology presence. They presented their findings at an important Joint Army–Navy–NASA–Air Force propulsion conference in Tennessee last June. Burton has also made presentations to DARPA and DOE. He says, "We now have a family of computational tools to use on sufficiently powerful computers and the experience implementing them, allowing us to study for the first time the performance of rocket engines and launch vehicles to unprecedented levels of detail."

House adds, "We have the HPC systems, codes, expertise, and additive-manufacturing tools to jump-start the rocket industry in a new direction." Addis predicts the adoption of HPC will spark cost-effective rocket engines and advanced launch vehicles that could eventually reduce launch costs from the current $10,000 per pound to about $100 per pound. He says, "We're on the brink of a rocket-engine revolution."

—Arnie Heller

Key Words: additive manufacturing, aerospike, Cab supercomputer, CharlesX, computational fluid dynamics (CFD), Defense Advanced Research Projects Agency (DARPA), Department of Defense (DOD), Flamelet Progress Variable (FPV), high-performance computing (HPC), JOE, large-eddy simulation (LES), liquid-propellant rocket engine (LRE), Next-Generation Rocket (NGR) program, Predictive Science Academic Alliance Program, Reynolds-Averaged Navier Stokes (RANS), Sequoia supercomputer, Stanford University Center for Turbulence Research (CTR), Syrah supercomputer, Turbulence Analysis and Simulation Center (TASC), XS-1 launch vehicle.

For further information contact Greg Burton (925) 422-3515 (burton29@llnl.gov) or Bill Bruner (925) 422-0552 (bruner6@llnl.gov).